

Evolution of MUMPS towards multicore architectures

Wissam Sid-Lakhdar

Ecole Normale Supérieure de Lyon, France

Abstract

We present some challenges of using multicore architectures and show adaptations of MUMPS to better benefit from these architectures. We show how to efficiently apply coarse/fine grain parallelisms by combining OpenMP and multi-threaded BLAS in the multifrontal factorization. Then, we present modifications to the internal dense factorization kernels to take into account memory contention and cache misses in a multi-threaded environment and lower computation times. Also, memory affinity issues on NUMA architectures are discussed. Finally, we show the impact of local scheduling policies (inside dense kernels) on hybrid shared-distributed memory architectures and propose techniques to scale on ever growing problem sizes. Experimental results are presented to validate the effectiveness of every algorithm and will give an idea of the performance that could be expected from MUMPS in the future.